Density-Based Region Search with Arbitrary Shape for Object Localization

Ji Zhao[1,*], Deyu Meng[2], Jiayi Ma[3]
1 Samsung Advanced Institute of Technology, Beijing 100027, China
2 School of Mathematics and Statistics, Xi'an Jiaotong University, Xi'an 710049, China
3 Electronic Information School, Wuhan University, Wuhan 430072, China
$*$ **E-mail: zhaoji84@gmail.com**

arXiv:1410.6447v1 [cs.CV] 23 Oct 2014

Abstract

Region search is widely used for object localization. Typically, the region search methods project the score of a classifier into an image plane, and then search the region with the maximal score. The recently proposed region search methods, such as efficient subwindow search and efficient region search, are much more efficient than sliding window search. However, for some classifiers and tasks, the projected scores are nearly all positive, and hence maximizing the score of a region results in localizing nearly the entire images as objects, which is meaningless.

In this paper, we observe that the large scores are mainly concentrated on or around objects. Based on this observation, we propose a method, named level set maximum-weight connected subgraph (LS-MWCS), which localizes objects with arbitrary shapes by searching regions with the densest score rather than the maximal score. The region density can be controlled by a parameter flexibly. And we prove an important property of the proposed LS-MWCS, which guarantees that the region with the densest score can be searched. Moreover, the LS-MWCS can be efficiently optimized by belief propagation. The method is evaluated on the problem of weakly-supervised object localization, and the quantitative results demonstrate the superiorities of our LS-MWCS compared to other state-of-the-art methods.

1 Introduction

In object localization, how to train the classifier and localize the objects is a chicken-and-egg problem [1]. Intuitively, if we know the locations of the objects, training a good classifier should be easier [2]; alternatively, if we have an ideal classifier, the higher scores of the classifier should distribute on or surrounding the objects, and localization should be easier. Therefore, it is common to consider these two tasks jointly. Given the spatial distribution of the classifier's score, finding the region that most probably containing object is known as region search problem [3, 4].

If we want to localize objects in a weakly-supervised manner, i.e., the label of a training image is in the image level, the dependency between classifier training and object region searching will be even stronger. The classifiers are typically trained by adopting the standard bag-of-words representation and support vector machines (SVMs), while the bottleneck is the region search. Here, region search aims to find candidate regions according to the spatial distribution of SVM scores. The popular sliding-window-based search, which exhaustively applies the classifier to rectangles within an image, is a natural way for region search. However, it is not efficient. Many heuristics have been used to speed up the search procedure, but they also introduce the risk of imprecisely localizing the object or even missing it. Many methods with global optimum and efficient solutions have been proposed to instead sliding window, here we name a few as representative ones.

The efficient subwindow search (ESS) is a remedy for sliding-window search [3]. ESS finds the rectangle that has the maximal score within an image. This turns out to be the 2D maximum subarray problem and the authors proposed an efficient branch-and-bound strategy to solve it. Inspired by the great success of ESS, many methods have been proposed to generalize the rectangle restriction in ESS. For example, the rectangle was generalized to composite boxes and polygons in [5]. The advantages of ESS lie in that it is parameter-free, and there exist efficient algorithms to solve it. It also has certain limitations. (i) If all projected scores are positive, ESS will find the entire

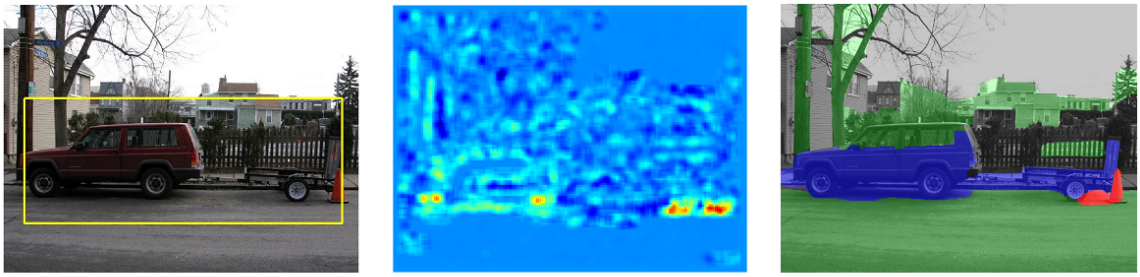

(a) (b) (c)

Figure 1. Detection by maximal score vs. detection by densest score. (a) ESS [3] that maximizes classifier's score tends to find much larger area than objects. (b) Visualization of score distribution for classifier (χ^2-SVM) [2]. (c) Localization of LS-MWCS. The green region denotes the detection with level 0 (corresponds to ERS [4]). By increasing the level, the localization shrinks from the green region to blue and red regions. Best viewed in color.

image to be the best subwindow. (ii) ESS can only find rectangular bounding boxes. For objects of non-rectangle shapes, noises in windows may mislead the localization, causing unsatisfactory results.

In order to find regions with arbitrary shapes, efficient region search (ERS) [4] was proposed for object localization. ERS converts region search into the maximum-weight connected subgraph (MWCS) problem, and it can be equivalently transformed into an instance of the prize-collecting Steiner tree (PCST) problem [6]. PCST is an NP-Complete problem but has efficient approximation solutions [7, 8]. The localization ability of ERS has been demonstrated in supervised classification framework. In [9], an adaptive image grids scheme was proposed to improve ERS by finding a better segmentation. ERS is successful for supervised localization. However, if all projected scores are positive, it will also find the entire image to be the best region as that in ESS.

Besides ESS and ERS, some region search methods are proposed for interactive segmentation. For example, twisted window search (TWS) [10] and shape from point features (SFF) [11] are proposed for interactive shape localization. SFF use α-complex [12] to build a filtration of simplicial complexes from a user-provided set of features. The densest complex connected component is viewed as the shape. SFF is successful for interactive shape localization. However, it relies on the distinct weights manually labeled. Empirically, it is not suitable for object localization because the projected weights from the classifier are usually very noisy compared to human interaction.

The weakly-supervised object localization has higher demand for region search than supervised localization, since the distribution of projected scores is in general very noisy. The basic idea of our method is based on the following observation: although the spatial distribution of a classifier's scores seems to be random, the densest area of the score distribution is on or near the object, see Fig. 1(b) for example. Instead of finding the maximal score as done in ESS and ERS, we attempt to search the regions with dense scores. The comparison of localization by densest score and maximal score is demonstrated in Fig. 1. Our region search method is called **LS-MWCS (level set maximum-weight connected subgraph)**, which has the following two advantages: (i) it is able to localize regions of arbitrary shape; (ii) it works well when all the weights are positive.

The contribution of this paper is two-fold. First, we propose a region search method that can be used for very noisy weight distributions, such as for weakly-supervised localization. Second, we validate the effectiveness of the densest score for region search, which in previous literatures is dominantly achieved by maximizing the score.

1.1 Related Work

Region search is a key technology in weakly supervised localization (WSL). WSL is usually modeled as multiple instance learning (MIL). In the MIL setting, each image is modeled as a bag of regions, and each region is an

instance. With two classes, the negative bag only contains negative instances and the positive consists of at least one positive. The goal of MIL is to label the positive instances in the positive bags. Region search corresponds to finding the region (instance) in the positive image (bag) that triggers the positive label. In the past few years, many MIL algorithms have been successfully used for weakly supervised learning, such as MILboost [13] and MI-SVM [1]. In [14], a region weighting method is proposed for WSL, which is customized for bag-of-words feature representation and non-linear SVM classifiers. Region search has a close relationship with common pattern discovery from images that share common contents, such as co-segmentation and image feature matching [15–17].

2 Density-Based Region Search

There are two schemes for region search: one is based on feature points [3], and the other is based on super-pixels from over-segmentation [4]. For methods based on feature points, it is difficult to localize regions that are consistent with the object boundaries. By contrast, localization methods based on super-pixels can guarantee the local consistency. In this paper, we focus on region search with super-pixels.

The effectiveness of ERS has been validated for supervised localization, while its performance for weakly-supervised localization has not been reported. Empirically, we found that ERS has the following problem when it is directly used in weakly-supervised localization. The shape and area of the connected component found by ERS cannot be directly controlled. In weakly supervised localization, since the positive weights are scattered over the entire image, ERS tends to find spindly shapes that cover nearly the whole image. However, we find that the densest area of the score distribution is on or near the object, see Fig. 1. Based on this observation, in this section we will extend ERS to find regions with densest scores.

Fist, we give some formal definition. The essence of ERS is the maximum-weight connected subgraph (MWCS) problem, which is defined as follows.

Definition 1 (Maximum-Weight Connected Subgraph (MWCS) Problem [4, 6]). *Given a connected undirected, vertex-weighted graph $G = (V, E)$ with weights $\omega : V \to \mathbb{R}$, find a connected subgraph $T = (V_T \subseteq V, E_T \subseteq E)$ of G to maximize the score $W(T) = \sum_{v \in V_T} \omega(v)$.*

We also define the density of a graph. In this paper, density means the average weight of a graph.

Definition 2 (Density of Graph). *Given a connected undirected, vertex-weighted graph $G = (V, E)$ with weights $\omega : V \to \mathbb{R}$, the density is defined as $\frac{\sum_{v \in V} \omega(v)}{|V|}$, where $|\cdot|$ means the node number of a graph.*

In the region search problem, the set of vertices V consists of the superpixels, and an edge in E connects a pair of superpixels that share a boundary. The weight of a vertex $\omega(V)$ is the superpixel's classifier score. The basic idea of ERS [4] is that the subgraph with maximal score corresponds to the localized object.

Since the MWCS problem aims to find a subgraph with maximum-weight, the density of the resulting subgraph cannot be controlled. In order to control the density of a subgraph, we define the level set MWCS (LS-MWCS) based on the definition of MWCS. As will be shown in Theorem 1, it has a close relationship with the density of the resulting subgraph.

Definition 3 (Level Set Maximum-Weight Connected Subgraph (LS-MWCS) Problem). *Given a connected undirected, vertex-weighted graph $G = (V, E)$ with weights $\omega : V \to \mathbb{R}$, the MWCS with level α aims to find a connected subgraph $T = (V_T \subseteq V, E_T \subseteq E)$ of G, that maximizes the score $W(T; \alpha) = \sum_{v \in V_T} (\omega(v) - \alpha)$.*

Based on the definition of LS-MWCS problem, we find the relationship between level α and the density of resulting subgraph. This property turns out to be very useful for density-based region search.

Theorem 1. *Denote the MWCS of a graph with level α_1 and α_2 as V_T^1 and V_T^2 respectively. Suppose $0 \leq \alpha_1 < \alpha_2$, then (i) the node number satisfies $|V_T^1| \geq |V_T^2|$; (ii) the weight summation satisfies $\sum_{v \in V_T^1} \omega(v) \geq \sum_{v \in V_T^2} \omega(v)$; (iii) if $|V_T^2| > 0$, then graph density satisfies $\frac{\sum_{v \in V_T^1} \omega(v)}{|V_T^1|} \leq \frac{\sum_{v \in V_T^2} \omega(v)}{|V_T^2|}$.*

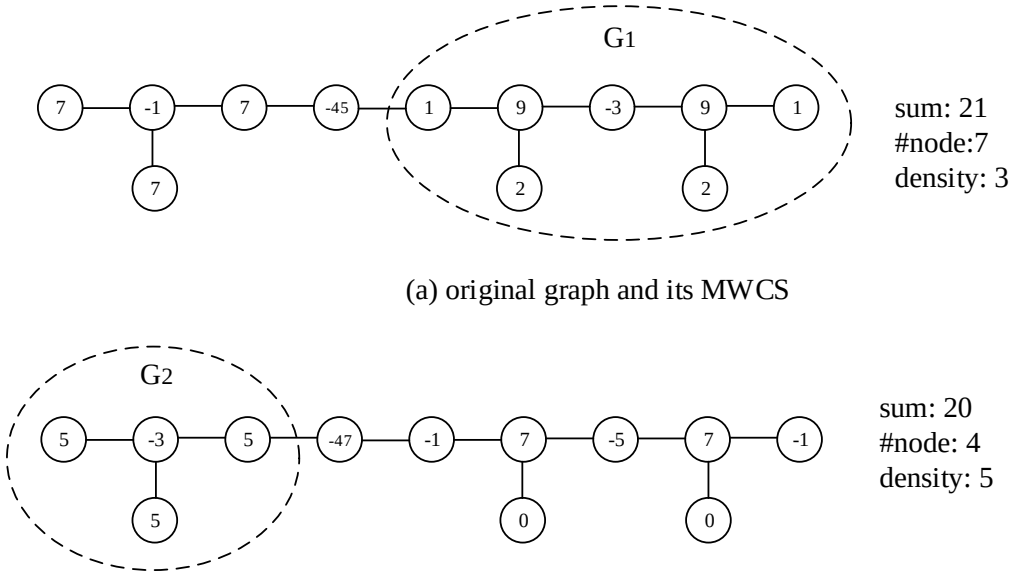

(a) original graph and its MWCS

(b) LS-MWCS with level = 2

Figure 2. An example of LS-MWCS.

Proof. (i) Since V_T^1 and V_T^2 are the MWCS with level α_1 and α_2, by Definition 3 we have

$$\sum_{v \in V_T^1} (\omega(v) - \alpha_1) - \sum_{v \in V_T^2} (\omega(v) - \alpha_1) \geq 0; \tag{1}$$

$$\sum_{v \in V_T^2} (\omega(v) - \alpha_2) - \sum_{v \in V_T^1} (\omega(v) - \alpha_2) \geq 0. \tag{2}$$

By adding the above two inequalities, we get $(\alpha_2 - \alpha_1)(|V_T^1| - |V_T^2|) \geq 0$. Since $\alpha_2 > \alpha_1$, the inequality $|V_T^1| \geq |V_T^2|$ holds.

(ii) Reformulate inequality (1) as $\sum_{v \in V_T^1} \omega(v) - \sum_{v \in V_T^2} \omega(v) \geq \alpha_1(|V_T^1| - |V_T^2|)$. Since we have proved that $|V_T^1| \geq |V_T^2|$, it is obvious that $\sum_{v \in V_T^1} \omega(v) \geq \sum_{v \in V_T^2} \omega(v)$ holds.

(iii) According to inequality (2), we have

$$\sum_{v \in V_T^2} \omega(v) - \alpha_2 |V_T^2| \geq \sum_{v \in V_T^1} \omega(v) - \alpha_2 |V_T^1|. \tag{3}$$

We have proved that $|V_T^1| \geq |V_T^2|$ and we suppose $V_T^2 > 0$, so the inequality still holds when we divide the left expression with $|V_T^2|$ and divide the right expression with $|V_T^1|$. □

According to Theorem 1, when we increase the level α, the node number and weight summation of the resulting subgraph is monotonically decreasing and the density of the subgraph is monotonically increasing. Fig. 2 gives an example. Fig. 2(a) is a graph with 12 vertices and 11 edges, the weights are labeled in corresponding nodes. According to Definition 1, it is easy to verify that its MWCS is subgraph G_1, where the node number, weight summation and density are 7, 21 and 3 respectively. If the weights of all the nodes minus 2, we get a new graph as

shown in Fig. 2(b). The MWCS of this new graph is subgraph G_2 . According to Definition 3, the LS-MWCS with level 2 is constructed as following: the node and edge are the same as the subgraph G_2 and the corresponding node weights are derived from the original graph. So the node number, weight summation and density of LS-MWCS are 4, 20 and 5 respectively.

Theorem 1 provides a theoretical background for finding dense subgraph. An interesting discovery in Theorem 1 is that the weight summation of the generated subgraph is monotonically decreasing with the level α. This is useful in practice. For example, if we want to find the dense subgraph with predefined weight summation, we can use the binary search to find the proper level α.

When the level α in LS-MWCS is determined, this problem can be easily converted to the standard MWCS problem. The MWCS problem is NP-complete. Fortunately, it can be transformed into an instance of the prize-collecting Steiner tree (PCST) problem, as shown in [6], which can be efficiently solved in practice for our problem setting ($10^1 \sim 10^2$ nodes, $10^2 \sim 10^3$ edges). In this paper we adopt an efficient belief propagation scheme to solve this problem[1] [8].

2.1 Parameter Setting for LS-MWCS

There is only one parameter, the level α, for the problem of LS-MWCS. The level controls the density of the generated subgraph. By increasing α, we can obtain the subgraph which has denser weights and vice versa. For LS-MWCS based region search in the paper, this parameter is very useful when the projected classifier scores are all positive, such as Hough transform vote [18] and matrix completion multi-label classifiers [19].

Generally, the selection of level α is task-dependent. Here we provide three strategies for the selection of α. First, the level α can be determined by a predefined node number, weight summation or graph density. Second, we can find an optimal level that has best performance for a specific task that depends on level α. Take region search for weakly-supervised localization as an example, the optimal level α can be set as the number producing the highest image classification accuracy by cross-validation. Third, we can find α that corresponds to stable subgraphs. Here the notation of stable is similar to that in the definition of maximally stable extremal regions (MSER) [15]. A connect subgraph is stable if its weight summation does not change too much when the level α is slightly changed. The efficient solution of stable subgraph is based on the fact that the weight summation is a monotonically decreasing function with level α, as proved in Theorem 1.

3 Experimental Results

In this section, we present the experiments of our density-based region search, including experiments on synthetic data and image datasets.

3.1 Results on Synthesis Data

Fig. 3 demonstrates an example of localizing regions using LS-MWCS. The image is 19×19 regular grids. The nodes of the graph are the 361 grids; the adjacency of the graph is determined by 4-connected neighborhood in the image plane; the weight of each node is the blueness of each block. The weights of the nodes are 2D Gaussian distribution with noise. It means that all the nodes have positive weights. The solution of MWCS will find all the nodes as subgraph, as shown in Fig. 3(a). By increasing the level α in LS-MWCS, the optimal subgraph shrinks to denser subgraphs, see Fig. 3(b) - (e).

3.2 Object Localization on Image Dataset

Next, we demonstrate the effectiveness of our method for weakly-supervised object localization. Similar to [1], we use a scheme of iterated optimization for classifier construction and region search.

[1] We use the code provided by the authors which is available from http://areeweb.polito.it/ricerca/cmp/code.

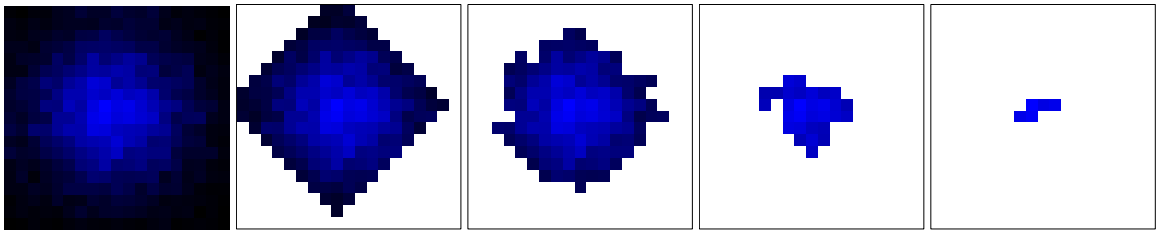

(a) $\alpha = 0$ (b) $\alpha = 0.1$ (c) $\alpha = 0.3$ (d) $\alpha = 0.7$ (e) $\alpha = 0.9$

Figure 3. Localization by MWCS with different level α. The image is 19×19 regular grids. The nodes of the graph are the grids; the adjacency of the graph is determined by 4-connected neighborhood in image plane; the weight of each node is the blueness of each block.

The dataset used in this experiment is Pittsburgh Car dataset [1], which contains 400 images, including 200 positive samples and 200 negative samples. There is only one car in each positive sample. Half of the positive and negative samples are used as training data, and the rest are used for testing.

The bag-of-words implementation is provided by VLFEAT toolbox [20]. It extracts dense SIFT feature uniformly [21], and generates the visual words by vector quantization. In all experiments, we extract 128 dimensional SIFT descriptors over a grid with a step of 5 pixels on fixed scale (16×16 pixels). The dictionary of 1000 visual words is obtained by clustering $100,000$ samples via K-means clustering. Linear SVM is adopted as classifier. The SVM parameter and the level α in LS-MWCS are chosen via cross-validation. The optimal α is the one that maximizes the image classification accuracy.

Fig. 4 demonstrates the qualitative performance comparison of 6 methods on Pittsburgh Car dataset. We consider methods that only use the scores of features. The 1st row in Fig. 4 is the visualization of feature weights[2]. The 2nd to 5th rows are results of ESS [3], TWS [10], SFF [11] and ERS [4], respectively. The 6th row is the results of our LS-MWCS. Clearly, our results are more visually pleasurable.

From Fig. 4, we can see that ESS is biased toward choosing the entire image as the object of interest, and it is not capable of localizing objects accurately. In previous literatures, the effectiveness of TWS and SFF are validated in interactive shape localization, where the features are pixel-level and the weight of feature is discretized to only 3 levels. Although TWS and SFF have potential applications in object localization, our experiments demonstrate that they cannot be used for such application directly and some modifications are needed.

To provide a quantitative measure for the localization performance, we compared methods using precision-recall curves, as shown in Fig. 5. We used the area of overlap (AO) measure to evaluate the correctness of localization. For this criterion, a threshold t should be defined for AO to imply a correct detection. In this paper, we set the threshold t to 0.4. The precision-recall curves are shown in Fig. 5. We compared our density-based region search method with other region search methods and weakly-supervised localization methods, including ESS [3], ERS [4], MILboost [13] and MI-SVM [1]. The average precisions (AP) of these methods are given in the legends. We can see that our LS-MWCS method performs significantly better than the baselines.

In our density-based region search, parameter α is the only parameter and is very important. If we set α as 0, our density-based region search is the same as ERS. From Fig. 4 and 5, we can see that our region search with level α performs better than original ERS in our experiment. This demonstrates that the control of score density by adjusting level α is effective for weakly-supervised localization.

[2]We blur the weight using a Gaussian filter and then use the visualization code of O. Woodford available at http://www.robots.ox.ac.uk/∼ojw/software.htm.

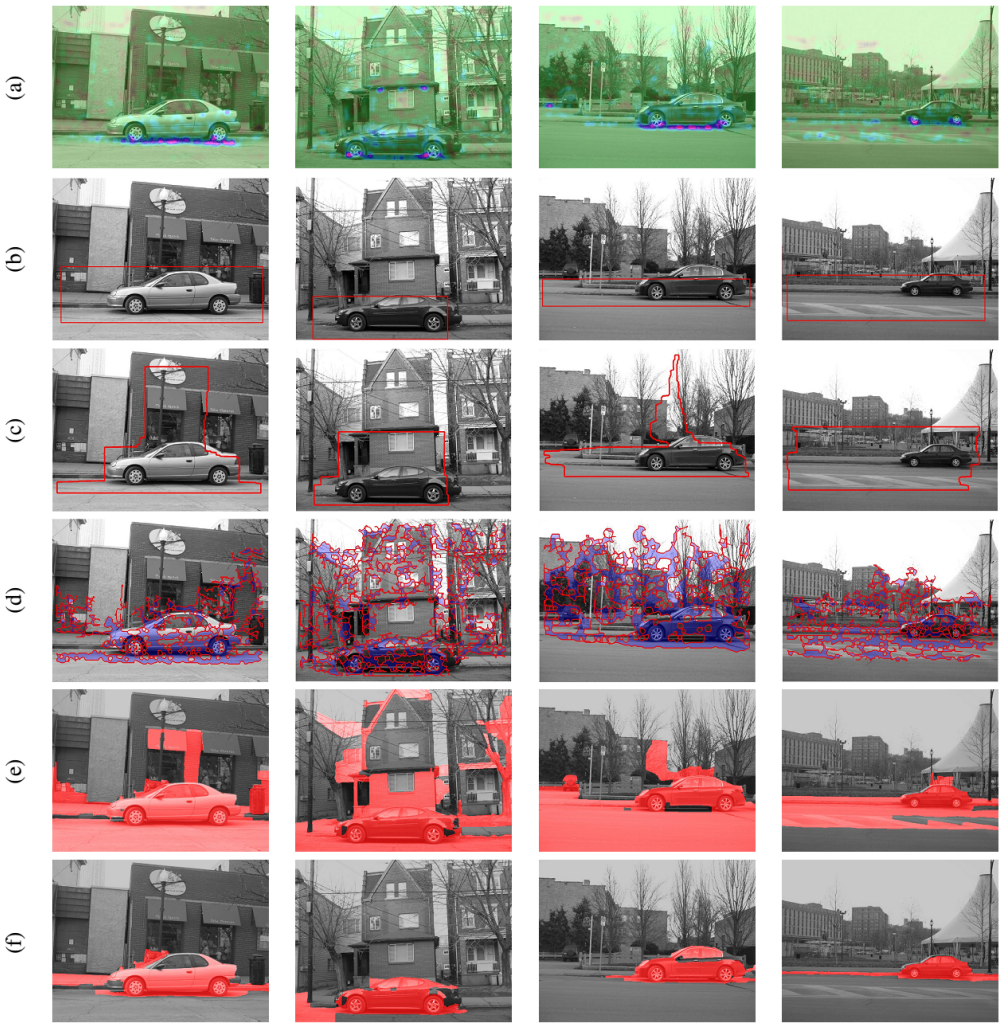

Figure 4. Performance comparison of region search methods. (a) Visualization of BoW; (b) ESS; (c) TWS with parameter $\lambda = 10^{-3}$; (d) SFF; (e) ERS; (f) LS-WMCS.

4 Conclusion

In this paper, we studied the problem of region search for object localization. We proposed a new strategy of finding the regions with densest score rather than the conventional maximal score. Based on this strategy, a density-based region search method was then developed, which is called LS-MWCS (level set maximum-weight connected subgraph). The effectiveness of our method was validated by synthesis data and weakly-supervised localization for images. Our method can be easily extended to videos.

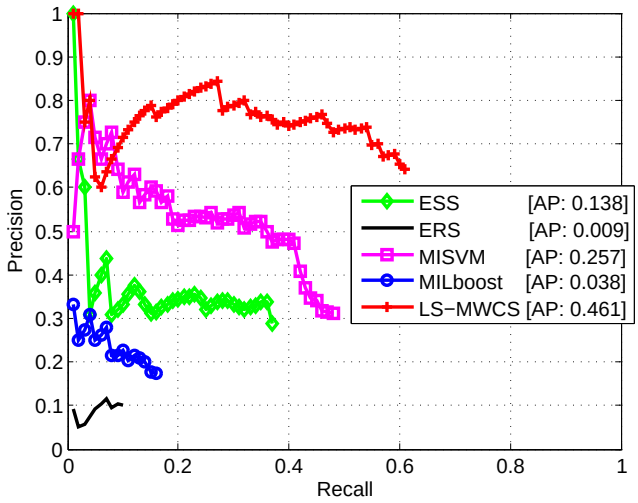

Figure 5. Localization performance comparison on Pittsburgh Car dataset.

Acknowledgments

We are very grateful to Fernando De la Torre at Carnegie Mellon University for helpful discussions. We thank Minh Hoai at The University of Oxford for providing the Pittsburgh Car dataset.

References

1. Nguyen MH, Torresani L, De la Torre F, Rother C (2009) Weakly supervised discriminative localization and classification: A joint learning process. In: Proceedings of International Conference on Computer Vision.

2. Uijlings JRR, Smeulders AWM, Scha RJH (2012) The visual extent of an object suppose we know the object locations. International Journal of Computer Vision 96: 46–63.

3. Lampert CH, Blaschko MB, Hofmann T (2008) Beyond sliding windows: Object localization by efficient subwindow search. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.

4. Vijayanarasimhan S, Grauman K (2011) Efficient region search for object detection. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.

5. Yeh T, Lee JJ, Darrell T (2009) Fast concurrent object localization and recognition. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.

6. Dittrich MT, Klau GW, Rosenwald A, Dandekar T, Müller T (2008) Identifying functional modules in protein-protein interaction networks: an integrated exact approach. Bioinformatics 24: i223–i231.

7. Ideker T, Ozier O, Schwikowski B, Siegel AF (2002) Discovering regulatory and signalling circuits in molecular interaction networks. Bioinformatics 18: S233–S240.

8. Bailly-Bechet M, Borgs C, Braunstein A, Chayes J, Dagkessamanskaia A, et al. (2011) Finding undetected protein associations in cell signaling by belief propagation. Proceedings of the National Academy of Sciences 108: 882–887.

9. Zhou C, Yuan J (2012) Arbitrary-shape object localization using adaptive image grids. In: Proceedings of Asian Conference of Computer Vision.

10. Gu S, Zheng Y, Tomasi C (2012) Twisted window search for efficient shape localization. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.

11. Gu S, Zheng Y, Tomasi C (2012) Shape from point features. In: Proceedings of International Conference on Acoustics, Speech and Signal Processing.

12. Edelsbrunner H, Kirkpatrick DG, Seidel R (1983) On the shape of a set of points in the plane. IEEE Transactions on Information Theory 29: 551–559.

13. Galleguillos C, Babenko B, Rabinovich A, Belongie S (2008) Weakly supervised object localization with stable segmentations. In: Proceedings of European Conference on Computer Vision.

14. Zhao J, Wang L, Cabral R, De la Torre F (2014) Feature and region selection for visual learning. arXiv preprint arXiv:1407.5245 .

15. Matas J, Chum O, Urban M, Pajdla T (2012) Robust wide baseline stereo from maximally stable extremal regions. In: Proceedings of British Machine Vision Conference.

16. Ma J, Zhao J, Tian J, Yuille AL, Tu Z (2014) Robust point matching via vector field consensus. IEEE Transactions Image Processing 23: 1706–1721.

17. Ma J, Chen J, Ming D, Tian J (2014) A mixture model for robust point matching under multi-layer motion. PloS ONE 9: e92282.

18. Maji S, Malik J (2009) Object detection using a max-margin Hough transform. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition.

19. Cabral RS, De la Torre F, Costeira JP, Bernardino A (2014) Matrix completion for weakly-supervised multi-label image classification. IEEE Transactions Pattern Analysis and Machine Intelligence .

20. Vedaldi A, Fulkerson B (2010) VLFeat - An open and portable library of computer vision algorithms. In: Proceedings of ACM Multimedia.

21. Lowe D (2004) Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60: 91–110.